(2)

UMIACS-TR-90-117
CS-TR-2531

September 1990

Combining Knowledge Bases Consisting of First Order Theories

Chitta Baral, Sarit Kraus, Jack Minker, and
V.S. Subrahmanian

Institute for Advanced Computer Studies and
Department of Computer Science
University of Maryland
College Park, MD 20742, U.S. A.

Abstract

Consider the construction of an expert system by encoding the knowledge of different experts. Suppose the knowledge provided by each expert is encoded into a knowledge base. Then the process of *combining* the knowledge of these different experts is an important and non-trivial problem. We study this problem here when the expert systems are considered to be first order theories. We present techniques for resolving inconsistencies in such knowledge bases. We also provide algorithms for implementing these techniques.

MASTER COPY KEEP THIS COPY FOR REPRODUCTION PURPOSES

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	Sep 90	Technical

4. TITLE AND SUBTITLE

Combining Knowledge Bases Consisting of First Order Theories

5. FUNDING NUMBERS

DAAG29-85-K-0177

6. AUTHOR(S)

Chitta Baral, Sarit Kraus, Jack Minker, V. S. Subrahmanian

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Maryland
College Park, MD 20742

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U. S. Army Research Office
P. O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

ARO 22648.20-EL

11. SUPPLEMENTARY NOTES

The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Consider the construction of an expert system by encoding the knowledge of different experts. Suppose the knowledge provided by each expert is encoded into a knowledge base. Then the process of *combining* the knowledge of these different experts is an important and non-trivial problem. We study this problem here when the expert systems are considered to be first order theories. We present techniques for resolving inconsistencies in such knowledge bases. We also provide algorithms for implementing these techniques.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

1 Introduction

Consider the construction of an expert system by encoding the knowledge of different experts. Suppose the knowledge provided by each expert is encoded into a knowledge base. Then the process of *combining* the knowledge of these different experts is an important and non-trivial problem.

The problem is important because the user of the expert system so constructed should have access to the knowledge of each of the experts. In particular, he/she should be able to to use the knowledge of two different experts to jointly derive a fact that neither of the experts, individually, knew. In other words, one important feature involved in consulting multiple experts is to *pool* their knowledge together and thus obtain knowledge that no individual expert previously had.

The problem is non-trivial because individual experts can, and often do, hold conflicting views on their domain of expertise. Two attorneys involved in a legal defense may well hold conflicting views on the best possible defense strategy, just as two doctors may well differ in their assessment of a patient's malady. In a logic knowledge base, these conflicting opinions manifest themselves in the form of inconsistencies. Classical logic would then indicate that the resulting knowledge base is meaningless – a state of affairs that is clearly inappropriate in this context. Just as the attorneys and doctors would work together to reconcile their views in the interests of the defendent or patient, so should a knowledge base management system reconcile these inconsistencies and allow sensible decisions to be drawn. The key problem here is: how should these inconsistencies be reconciled? This is the problem addressed in this paper.

Baral, Kraus and Minker [BKM89] formalize the notion of combining knowledge bases when each knowledge base is a general Horn logic program (set of rules with only atoms allowed in the head) and assume that the union of the knowledge bases is stratified (no recursion through negation). They assume the presence of world knowledge in the form of integrity constraints, which all the individual knowledge bases satisfy, and the combined knowledge base is required to satisfy. They present methods to obtain a maximally combined knowledge base with respect to the union of knowledge bases that is consistent with respect to the integrity constraints. Since they consider each knowledge base to be a general Horn logic program the union of the knowledge bases is always consistent and the combined knowledge base they obtain is maximal with respect to it.

In this paper we consider each knowledge base as a first order theory. The set of integrity constraints is also assumed to be a first order theory. In this case the union of the knowledge bases is not necessarily consistent. Because of this, in the absence of integrity constraints, to have the combined knowledge base as the union of the knowledge bases, we need a semantics for inconsistent theories. Many such semantics have been suggested in the past [BS89, BS88, dCSV89, Gra74, Gra75, Gra77, Gra78, Sub89, GL, KS]. In this paper we use the *cautious* approach of Grant and Subrahmanian [GS90] to characterize the semantics of inconsistent theories. In the cautious approach the semantics of an inconsistent theory is the semantics obtained by considering all maximally consistent subsets of the inconsistent theory. A sentence is considered true (false) if it is true (false) in all maximally consistent subsets of the original theory.

In the next section, we will present a scenario which will be used throughout the paper to illustrate the basic intuitions behind our technical development. In section 3 of the paper, we discuss the cautious semantics of inconsistent theories in the presence of integrity constraints. In the subsequent section we formalize combining a set of theories having the same priority, in the presence of integrity constraints and its relationship with view update approaches [FKUV86, FUV83]. We then allow the theories to be prioritized and formalize the notion of combining a set of prioritized theories.

2 A Motivating Scenario

Inconsistencies can easily arise when *multiple* reasoning agents each arrive at a particular view of the world. While these individual views are usually self-consistent, they often tend to conflict with one another. We now present a simple scenario which we will use over and over again to motivate the basic ideas in the paper.

The Scenario: At 1:00 AM on January 14, 1990, Don was shot outside the Good Times Bar in Washington. The street was more or less deserted (it being late in the night) except for four people: Don (who got shot), the murderer, and two rather drunk individuals, John and Bill, who were on the street. John is an eighty-five year old man who was about 100 yards away from the shooting, while Bill is thirty years old. Bill was about 75 yards away from the shooting. Their stories are the following:

John's Story:

1. The murderer wore an orange coat.
2. The murderer wore no hat.
3. He knows the murderer got away in a car (as he heard the engine revving up and the car taking off), but he was hiding in a doorway and was too scared to look, and hence cannot tell us anything about the car.

Bill's Story:

1. The murderer wore a dark (probably black) coat.
2. The murderer wore no hat.
3. The murderer drove off in a pink Mercedes.

If we look at John's story and Bill's story, they are self-consistent. If Bill had not been around, we would probably have accepted John's version of the story (and vice-versa). However, their stories conflict with each other (if we make the reasonable assumption that the murderer wore only one coat). This assumption has the status of an integrity constraint: for the purposes of the story, it is a statement all parties are willing to accept.

Integrity Constraints:

1. The murderer wore only one coat at the time of the murder.
2. Based on other evidence, the police present a convincing case that the murderer knew the victim well.
3. Don's close cronies are Jeff, Ed and Tom.
4. There is no evidence that any of these three individuals had either borrowed or bought a coat recently; so the only coats they could have worn were their own.
5. Both Jeff and Ed have pink Mercedes. Tom doesn't know how to drive.
6. Jeff has an orange coat.
7. Ed has a black coat.
8. there is no possibility of any collusion between Jeff and Ed.

Based on the above story, we are led to suspect Ed or Jeff, but not both. Only one of them was the murderer. If we accept John's story, then Jeff is the murderer. If we accept Bill's story, then Ed is the murderer.

We are faced with the following problem: who did it ? There are numerous alternatives:

Alternative 1: In a court of law the guilt of a person must be established beyond ALL reasonable doubt. This cannot be established in this case. For example, if Jeff is on trial for the murder of Don, then a reasonable doubt can be cast on his guilt by the defense. A similar situation would occur if Ed was on trial. This situation corresponds to the case where one views each and every possibility in the correctness of the witnesses' statements. We accept a person as guilty iff he/she turns out to be guilty in in all these different possible worlds.

Alternative 2: We may be led to doubt the correctness of John's statements. After all, he is eighty-five years old, as compared to Bill's thirty years, and hence, presumably, Bill's eyesight is better. Furthermore, Bill was much closer to the scene of the crime, and hence, one may feel that his evidence is more credible. In deciding who to prosecute (Ed or Jeff), the police may well decide that they can make a more compelling case against Ed based on Bill's evidence. This alternative corresponds to the assignment of a *priority* to Bill's evidence, rather than to John.

Alternative 3: The third alternative is to simply conclude that the evidence is inconclusive: we know for sure that either Jeff or Ed did it, but cannot figure out which of them was actually responsible. This may trigger a search for further evidence which may, perhaps, yield something moe conclusive.

These are only three possible scenarios. Each of these represents a reasonable way of reasoning about the body of evidence in front of us. We will illustrate the technical development of the paper by frequent reference to this example.

3 Cautious Semantics for Inconsistent Theories

We consider a theory to be a finite set of well-formed formulas. Several semantics for inconsistent theories have been discussed in [GS90]. One of the approaches to characterize an inconsistent theory is to consider the maximally consistent subsets of the inconsistent theory. A maximally consistent subset of an inconsistent theory T

is a theory which is a consistent subset of T, and which becomes inconsistent if any other sentence of T is added to it. Intuitively, each maximally consistent subset of T corresponds to a consistent state of the world T is trying to characterize. In the presence of a set of worlds, one can either be bold and pick one of them as the "real" state of the world or one can be cautious (some call it skeptical) and consider all of them. Hence, in the cautious characterization of inconsistent theories, the truth value of a sentence L corresponds to the intuition: "Is L true w.r.t. each and every maximally consistent state of affairs ? ".

But, in the presence of integrity constraints (world knowledge which every theory has to satisfy) we have to consider only those worlds that are consistent with respect to the integrity constraints. The bold approach of doing it would be to consider only those maximally consistent subsets that agree with the integrity constraints. The cautious approach of doing it would be to subdivide a maximal consistent subset P of the original theory to maximal consistent subsets of P, such that they are each consistent with respect to the integrity constraints.

Example 3.1 With respect to the murder example described in the previous section, the cautious semantics would accept the conclusion "X is the murderer" iff X was the murderer irrespective of whether we chose to believe John or Bill. Thus, according to the cautious semantics, there would be no such individual X. However, the cautious semantics would allow us to conclude the sentence "Either Ed is the murderer or Jeff is the murderer" because this follows irrespective of whether we believe John or Bill.

In the following definitions $MAXCONS(P)$ and $MAXCONS(P, IC)$ are maximal consistent subsets of P and maximal consistent subsets of P with priority to IC. By, maximal consistent subsets of P with priority to IC we mean maximal consistent subsets of $P \cup IC$, which contain all elements of IC.

Definition 3.1 Let P be a theory and IC be a set of integrity constraints. A subset $Q \subseteq P \cup IC$ is said to be ***maximally consistent*** with priority to IC iff Q is consistent, $IC \subseteq Q$ and for every theory Q' such that $Q \subset Q' \subseteq P \cup IC$, it is the case that Q' is inconsistent. $MAXCONS(P, IC)$ is the set of maximally consistent subsets of $P \cup IC$ with priority to IC. When IC is an empty set then $MAXCONS(P, IC)$ is called $MAXCONS(P)$ and is the set of maximally consistent subsets of P.

Theorem 1 Suppose P is any first order theory and IC is any consistent set of integrity constraints. Then $P \cup IC$ has at least one maximally consistent subset P' such that $IC \subseteq P'$. (Note in particular, that P may contain function symbols and P may be infinite).

Proof. $P \cup IC$ has at least one consistent subset that is a superset of IC, viz. IC itself. Thus, let $CONS(P, IC)$ denote the set $\{X \mid X \subseteq P \cup IC$ and X is consistent and $IC \subseteq X\}$. Thus, $CONS(P, IC) \neq \emptyset$ because $IC \in CONS(P, IC)$. We show below that every ascending chain of elements in $CONS(P)$ has an upper bound in $CONS(P)$. The result then follows from Zorn's Lemma.

Suppose $M_1 \subseteq M_2 \subseteq M_3 \subseteq \cdots$ is an ascending sequence of members of $CONS(P)$, i.e. each M_i is a consistent subset of $P \cup IC$ and $IC \subseteq M_i$. Then $M = \bigcup_{i=1}^{\infty} M_i$ is an upper bound for this ascending sequence. Moreover, M is consistent, $IC \subseteq M$ and $M \subseteq P \cup IC$, i.e. $M \in CONS(P)$. The only non-obvious part is the consistency of M.

To see this, suppose M is not consistent. Then, by the Compactness Theorem, there is a finite subset $M' \subseteq M$ such that M' is inconsistent. Let $M' = \{\gamma_1, \ldots, \gamma_n\}$ for some integer n. Hence, for each $1 \leq i \leq n$, there is an integer, denoted α_i such that $\gamma_i \in M_{\alpha(i)}$. Let $\alpha = max\{\alpha(1) \ldots, \alpha(n)\}$. Then $M' \subseteq M_\alpha$. Hence, as M' is inconsistent, M_α is also inconsistent, thus contradicting our assumption that each M_j, $j \geq 1$, is in $CONS(P)$. □

A weaker version of the above theorem has been established by Grant and Subrahmanian [GS90] (cf. Corollary 3.1 below).

Corollary 3.1 [GS90] Every theory P has at least one maximally consistent subset.

Proof. Take IC to be the empty set in the proof of Theorem 1. □

Definition 3.2 [GS90] Suppose P is a theory, and F is a formula. A notion of entailment, $\vdash_\forall$ based on the cautious approach is defined as follows:
$P \vdash_\forall F$ iff $P' \models F$ for every maximal consistent subset $P' \subseteq P$.

Theorem 2 [GS90] Suppose P is a theory, and L, L_1, L_2 are ground literals. Then:

1. for all ground literals L, it is not the case that $P \vdash_\forall L$ and $P \vdash_\forall \neg L$.

2. $P \vdash_{\vee} (L_1 \,\&\, L_2)$ iff $P \vdash_{\vee} L_1$ and $P \vdash_{\vee} L_2$.

3. $P \vdash_{\vee} F$ for all tautologies F of classical logic. □

Example 3.2 Let P be:

$$p \vee \neg q$$

$$\neg p \vee \neg q$$

$$r$$

$$q$$

In this case, $MAXCONS(P)$ consists of three elements:

$\{p \vee \neg q\,;\ r\,;\ q\,\}$,
$\{\neg p \vee \neg q\,;\ r\,;\ q\,\}$,
$\{p \vee \neg q\,;\ \neg p \vee \neg q\,;\ r\,\}$.

In this case, $P \vdash_{\vee} r$. But $P \not\vdash_{\vee} q$ and $P \not\vdash_{\vee} p$ and $P \not\vdash_{\vee} \neg q$ and $P \not\vdash_{\vee} \neg p$.

We now present algorithms to compute $MAXCONS(P)$ and $MAXCONS(P, IC)$. The algorithms assume that the theories have a finite Herbrand Base. Function-free databases satisfy this condition. In other words it is decidable to determine their consistency. When an inconsistent theory T consists of n sentences, Algorithm 3.1 constructs its maximal consistent subsets by determining the consistency of each subset of T of cardinality $n-1$. All such consistent subsets are stored in a set S. For each inconsistent subset, its maximal consistent subsets are added to S. The set of maximal elements of S is $MAXCONS(P)$. Algorithm 3.2 is similar to Algorithm 3.1 except that instead of testing the consistency of each subset of T, it tests the consistency of the union of IC with, each of the subsets of T.

Algorithm 3.1 Procedure $MAXCONS1(P)$
$MAXCON = \emptyset$
If P *is consistent* **then** $MAXCONS1(P) = \{P\}$.
else
begin
{*** Let P be the set of sentences* $\{C_1, \cdots, C_n\}$. ****}
For $i = 1 \cdots n$ **do** $P_i := P - \{C_i\}$ **od**

For $i = 1 \cdots n$ **do** $MAXCON := MAXCON \cup MAXCONS1(P_i)$ **od**
$MAXCONS1(P)$:= *maximal elements of* $MAXCON$.
end

Theorem 3 $MAXCONS(P) = MAXCONS1(P)$

Proof: $[MAXCONS(P) \subseteq MAXCONS1(P)]$ Suppose $X \in MAXCONS(P)$. Then $X = P - \{C_1, \ldots, C_r\}$ for some integer $r \geq 0$ where $\{C_1, \ldots, C_r\} \subseteq P$. We proceed by induction on r.

Base Case. $(r = 0)$ In this case $MAXCONS(P) = \{P\} = MAXCONS1(P)$.

Inductive Case. $(r = k + 1)$ Consider $P' = P - \{C_{k+1}\}$. Then X is a maximal consistent subset of P' and furthermore, $X = P' - \{C_1, \ldots, C_k\}$. Therefore, by the induction hypothesis, $X \in MAXCONS1(P')$. As $X \cup \{C_{k+1}\}$ is inconsistent, X is in $MAXCONS1(P)$.

$[MAXCONS1(P) \subseteq MAXCONS(P)]$ Similar. □

Algorithm 3.2 Procedure $MAXCONS1(P, IC)$
$MAXCON = \emptyset$
If $P \cup IC$ *is consistent* **then** $MAXCONS1(P, IC) = \{P \cup IC\}$.
else
begin
{** *Let* P *be the set of sentences* $\{C_1, \cdots, C_n\}$. **}
For $i = 1 \cdots n$ **do** $P_i := P - \{C_i\}$ **od**
For $i = 1 \cdots n$ **do** $MAXCON := MAXCON \cup MAXCONS1(P_i, IC)$ **od**
$MAXCONS1(P, IC)$:= *maximal elements of* $MAXCON$.
end

Theorem 4 $MAXCONS(P, IC) = MAXCONS1(P, IC)$

Proof: Proceeds along exactly the same lines as the proof of Theorem 3. □

Returning to the motivating murder example in Section 2, if we take P to be the union of John's evidence and Bill's evidence, then $MAXCONS(P, IC)$ consists of two theories T_1 and T_2. T_1 contains:

1. all integrity constraints and

2. sentences 1,2, 3 of John's story and sentences 2, 3 of Bill's story.

Likewise, T_2 contains:

1. all integrity constraints and

2. sentences 2 and 3 of John's story and sentences 1,2, 3 of Bill's story.

Thus, using the $MAXCONS(P, IC)$ semantics, we may conclude that the murderer wore no hat (this being true in both T_1 and T_2 above). However, we may not conclude anything about the color of the murderer's coat. We may also conclude that the murderer drove away in a pink Mercedes.

We now discuss a technique for computing $MAXCONS(P)$ in cases when P is a finite set of clauses (a clause is a disjunction of literals). Throughout the rest of this section, we consider only sets of clauses.

Definition 3.3 Suppose T is a consistent set of clauses and D is a clause such that $T \cup \{D\}$ is inconsistent. A *refutation* of D from T is a sequence $C_1, \ldots, C_n$ such that:

1. C_n is the empty clause $\Box$ and

2. each C_i, $1 \leq i \leq n$ is either in $T \cup \{D\}$ or is a resolvent of two clauses $C, C' \in T \cup \{D\} \cup \{C_1, \ldots, C_{i-1}\}$.

$C_1, \ldots, C_n$ is called a *minimal* refutation of D from T if there is no strict subsequence of $C_1, \ldots, C_n$ which is also a refutation of D from T, i.e. there is no sequence $D_1, \ldots, D_m$ such that $\{D_1, \ldots, D_m\} \subset \{C_1, \ldots, C_n\}$ and if $D_i = C_j$ and $D_{i+1} = C_k$, then $j < k$.

Definition 3.4 Suppose T is a consistent set of clauses and D is a clause such that $T \cup \{D\}$ is inconsistent. Let $\Re$ be some minimal refutation of D from T. Then the set $\Re \cap T$ is said to be a *potential cause* of D.

Example 3.3 Suppose T is the following set of clauses:

$C1$: $a \vee b$

$C2$: $a \vee \neg b$

$C3$: $a \vee c$

$C4$: $a \vee \neg c$

and $D \equiv \neg a$. There are two minimal refutations $\Re_1$ and $\Re_2$ of D from T where:

$$\Re_1 = C1, C2, a, D, \Box$$

$$\Re_2 = C3, C4, a, D, \Box$$

(Actually a few more minimal refutations may be obtained by re-arranging the occurrences of some of the clauses in $\Re_1, \Re_2$). Thus the potential causes of D are: $PC_1 = \{C1, C2,\}$ and $PC_2 = \{C3, C4\}$ corresponding to $\Re_1, \Re_2$ respectively.

In the context of the murder scenario outlined in Section 2, there is only one potential cause of the inconsistency, viz. the following three sentences:

1. John: The murderer wore an orange coat.
2. Bill: The murderer wore a black coat.
3. IC: The murderer wore only one coat at the time of the murder.

If T is itself inconsistent, we may talk of refutations of the empty clause, $\Box$, from T. This refers simply to different refutations of $\Box$ from T. In this case, each minimal refutation of T gives rise to a *potential cause* of the inconsistency of T, viz. the potential cause of $\Box$ w.r.t. the refutation we are currently considering.

Algorithm 3.3 Procedure $MAXCONS2(P)$
Let there be n potential causes of the inconsistency of P.
If $n = 0$ **then** $MAXCONS2(P) = \{P\}$.
else
begin
For $i = 1 \cdots n$ **do** $S_i :=$ *the i'th potential cause of the inconsistency of P* **od**
$S := S_1 \times \cdots \times S_n$
$S' := \{\{a_1, \ldots, a_n\} \mid (a_1, \ldots, a_n) \in S\}$
$MINS :=$ *minimal elements of* S' *w.r.t inclusion*
$MAXCONS2(P) = \{P - Y, \mid Y \in MINS\}$
end

In the remaining Lemma and Theorems of this section, any unexplained notation will refer to the notation used in the $MAXCONS2$ algorithm.

Lemma 3.1 Suppose P is a set of clauses and $Z \in MINS$. Then $(P - Z)$ is consistent.

Proof. Suppose $(P - Z)$ is inconsistent. Then there is a minimal refutation of $(P - Z)$. Let $z_1, \ldots, z_m$ be the members of $(P - Z)$ (and hence of P) occurring in this refutation. Therefore, there exists an $1 \leq i \leq n$ such that $S_i = \{z_1, \ldots, z_m\}$. But then, as $Z \in MINS$, there must exist a $1 \leq j \leq m$ such that $z_j \in Z$. This implies that $z_j \notin (P - Z)$; a contradiction. □

Theorem 5 $MAXCONS(P) = MAXCONS2(P)$

Proof: $[MAXCONS(P) \subseteq MAXCONS2(P)]$ Suppose $X \in MAXCONS(P)$. It suffices to show that $(P - X) \in MINS$. To do this, we need to show two things.

I. First we show that $(P-X) \cap S_i \neq \emptyset$ for all $1 \leq i \leq n$. Suppose $(P-X) \cap S_i = \emptyset$ for some $1 \leq i \leq n$. Then $S_i \subseteq X$ which contradicts the assumption that X is consistent.

II. Next, we show that $(P - X)$ is a minimal element of S' w.r.t. inclusion. Suppose not. Then there is a $Z \in MINS$ such that $Z \subset (P - X)$. In particular, there exists an $a \in (P - X) - Z$. But then, by Lemma 3.1, $(P - Z)$ would be a consistent subset of P. But $X \subset (P - Z)$ thus contradicting the maximality of X.

$[MAXCONS2(P) \subseteq MAXCONS(P)]$ Suppose $X \in MAXCONS2(P)$. Let $S_1, \ldots, S_n$ be all the potential causes of the inconsistency of P. If $n = 0$, then $X = P \in MAXCONS2(P)$.

So assume $n > 0$. Then $X = (P - Y)$ where $Y = \{a_1, \ldots, a_n\}$ and for all $1 \leq i \leq n$, $a_i \in S_i$ and $Y \in MINS$.

Claim 1: X is consistent.
Proof of Claim 1: Suppose not. Then there exists an $1 \leq i \leq n$ such that $S_i \subseteq X$. In addition, $a_i \in S_i$. But $a_i \notin X$ because $X = (P - Y)$. But this contradicts the statement that $S_i \subseteq X$.

Claim 2: X is maximally consistent.
Proof of Claim 2: Suppose not, i.e. there exists a maximal consistent X' such that $X \subset X' \subseteq P$. let $X = (P - Y)$ and $X' = (P - Y')$. As $X \subset X'$, $Y' \subset Y$. As X' is a

maximal consistent subset of P, by the $[MAXCONS(P) \subseteq MAXCONS2(P)]$ part of the proof, $X' \in MAXCONS2(P)$. Thus, $Y' \in MINS$. But $Y \in MINS$ also. But this is a contradiction because $Y' \subset Y$. □

Algorithm 3.4 Procedure $MAXCONS2(P, IC)$
Let there be n potential causes of the inconsistency of $P \cup IC$
If $n = 0$ **then** $MAXCONS2(P, IC) = \{P \cup IC\}$.
else
begin
For $i = 1 \cdots n$ **do** $S_i :=$ *(the i'th potential cause of the inconsistency) -* IC **od**
$S := S_1 \times \cdots \times S_n$
$S' := \{\{a_1, \ldots, a_n\} \mid (a_1, \ldots, a_n) \in S\}$
$MINS :=$ *minimal elements of* S' *w.r.t. inclusion*
$MAXCONS2(P, IC) = \{(P - Y) \cup IC \mid Y \in MINS\}$
end

Theorem 6 $MAXCONS(P, IC) = MAXCONS2(P, IC)$

Proof: Similar to the proof of Theorem 5. □

4 Combining General Theories

The problem of combining general theories is formalized as follows. We have a set of consistent theories and a set of integrity constraints. Each theory satisfies the integrity constraints. We would like to combine the given set of theories so that the combined theory is also consistent with respect to the integrity constraints and contains as much consistent information as possible.

Definition 4.1 A theory is a finite set of sentences in first order logic. A flock (a term borrowed from [FKUV86]) is a set of theories, and the flock corresponding to an inconsistent theory T is the set of maximally consistent subsets of T.

Definition 4.2 Relation between flocks
Let F_1 and F_2 be flocks.
1. $\leq_1$: $F_1 \leq_1 F_2$ iff $(\forall T \in F_1)(\exists T' \in F_2) : T \subseteq T'$.
2. $\leq_2$: $F_1 \leq_2 F_2$ iff $(\forall T \in F_1)(\exists T' \in F_2) : Cn(T) \subseteq Cn(T')$, where $Cn(T)$ is the set of consequences of T.

Definition 4.3 Consistency

A flock F is said to be **consistent** with respect to a set of integrity constraints IC, iff for every theory T present in F, $T \cup IC$ is consistent. □

Definition 4.4 Correctness

A theory T is said to be **$\leq_1$-correct** with respect to theories $T_1, \cdots, T_k$, if $\{T\} \leq_1 \{T_1 \cup \cdots \cup T_k\}$. $\leq_2$-correctness is defined analogously. □

Definition 4.5 Combination of Theories

Let $T_1, \cdots, T_k$ be a set of theories and IC a set of integrity constraints; where each T_i satisfies IC. A combination function C is a mapping from a set of theories and a set of integrity constraints into a flock satisfying the following three criteria.

1.(identity) $C(\{T\}, IC) = T$.

2.(consistency) $C(\{T_1, \cdots, T_k\}, IC)$ is consistent with respect to IC.

3.(correctness) $C(\{T_1, \cdots, T_k\}, IC)$ is $\leq_1$-correct with respect to the theories $T_1, \cdots, T_k$.

Another useful property is associativity, which is defined as follows.

$C(\{T_1, \cdots, T_j, C(\{T_{j+1}, \cdots, T_k\}, IC)\}, IC)$

$=_{mm} C(\{T_{j+1}, \cdots, T_k, C(\{T_1, \cdots, T_j\}, IC)\}, IC)$

$=_{mm} C(\{T_1, \cdots, T_k\}, IC)$; where P $=_{mm}$ Q means $Cn(P) = Cn(Q)$. □

We now define three combination functions. The first two are skeptical in nature. The first combination function takes the union of the theories and the integrity constraints and looks at the maximal consistent subsets of this union with priority to the integrity constraints. More formally,

Definition 4.6 $Comb_1(\{T_1, \cdots, T_k\}, IC) \stackrel{\text{def}}{=} MAXCONS(T_1 \cup \cdots \cup T_k, IC)$.

Theorem 7 $Comb_1$ is a combination function, i.e. it satisfies the identity, consistency and correctness criteria. □

Note that the combination function $Comb_1$ takes the union of theories $T_1, \ldots, T_k$ and then finds the maximal subsets that are consistent with the integrity constraints,IC.

The second combination function takes the union of the theories and finds its maximal consistent subsets. It then looks at all the theories in $MAXCONS(Y_i, IC)$ for all $Y_i \in MAXCONS(T_1 \cup \cdots \cup T_k)$.

Definition 4.7 $Comb_2(\{T_1, \cdots, T_k\}, IC) \stackrel{\text{def}}{=}$ maximal elements of S, where $S = \{X : X \in MAXCONS(Y_i, IC)$ where $Y_i \in MAXCONS(T_1 \cup \cdots \cup T_k)\}$.

The following example illustrates the two approaches.

Example 4.1 Let the union of the theories T be:

a

b

c

$\neg a \vee \neg b$

and the set of integrity constraints be $IC = \{\neg a \;\vee\; \neg c\}$

$MAXCONS(T, IC) = \{\{a, b, \neg a \vee \neg c\}, \{a, \neg a \vee \neg c, \neg a \vee \neg b\}, \{b, c, \neg a \vee \neg c, \neg a \vee \neg b\}\}$

$MAXCONS(T) = \{T_1 = \{a, b, c\}, T_2 = \{a, c, \neg a \vee \neg b\}, T_3 = \{b, c, \neg a \vee \neg b\}\}$

$MAXCONS(T_1, IC) = \{\{a, b, \neg a \vee \neg c\}, \{c, b, \neg a \vee \neg c\}\}$

$MAXCONS(T_2, IC) = \{\{a, \neg a \vee \neg c, \neg a \vee \neg b\}, \{c, \neg a \vee \neg c, \neg a \vee \neg b\}\ \}$

$MAXCONS(T_2, IC) = \{\{b, c, \neg a \vee \neg c, \neg a \vee \neg b\}\}$

Hence, in this example $MAXCONS(T, IC) =$ maximal members of the set $\bigcup_{T_i \in MAXCONS(T)} MAXCONS(T_i, IC)$.

$a \vee b$ is true in all models of members of $MAXCONS(T, IC)$ and hence it is true with respect to the combination of these theories.

To see how $Comb_2$ behaves w.r.t. the murder scenario of Section 2, $Comb_2$ would first construct $MAXCONS(P)$ where P is the union of John's story and Bill's story. Note that P is perfectly consistent and hence $MAXCONS(P) = \{P\}$. (The fact that the murderer wore only one coat at the time of the murder is neccessary for the inconsistency to arise). $Comb_2$ now computes the maximal elements of $\{X \mid X \in MAXCONS(Y_i, IC)$ where $Y_i \in MAXCONS(P)\}$. In this case, this leads to exactly the same results as $Comb_1$. That these two seemingly different combination functions are indeed actually the same is now shown in the following theorem.

Theorem 8 $MAXCONS(T, IC) =$ maximal members of the set S:
where, $S = \{X : X \in MAXCONS(T_i, IC)\, where\, T_i \in MAXCONS(T)\}$

Proof:

$\Rightarrow$

Let $X \in MAXCONS(T, IC)$. Then there exists a Y in $MAXCONS(T)$ such that:

(1) $X \cap T \subseteq Y$ and (2) $X \in MAXCONS(Y, IC)$
Consider such a Y. Suppose $X \notin MAXCONS(Y, IC)$. Since, $X \cup IC$ is consistent, this means that there is an X' such that $X \subset X'$, and such that $X' \in MAXCONS(Y, IC)$, but then our assumption that $X \in MAXCONS(T, IC)$ is contradicted. Hence, $X \in MAXCONS(Y, IC)$.
This proves that $X \in S$. Since, $X \in MAXCONS(T, IC)$, it has to be a maximal member of S.

$\Leftarrow$

Suppose X is a maximal element of S. Then there is a $T_i \in MAXCONS(T)$ such that $X \in MAXCONS(T_i, IC)$ and Y such that $X \subset Y$ and $Y \in MAXCONS(T_j, IC)$ for some $T_j \in MAXCONS(T)$. This implies that X is consistent (note that by definition of $MAXCONS$, $IC \subseteq X$ and so X is consistent w.r.t. IC).

Suppose $X \notin MAXCONS(T, IC)$. Then there is an α such that $X \subset \alpha$ and $\alpha \in MAXCONS(T, IC)$. By the first part of the theorem, this means, that $\alpha \in MAXCONS(T_j, IC)$, for some $T_j \in MAXCONS(T)$. This violates our initial assumption about X, which says that no such T_j exists. A contradiction. Hence, $X \in MAXCONS(T, IC)$. □

Corollary 4.1 $Comb_2$ is a combination function. □

The third function uses a bold approach. Here we consider any maximal consistent subset of the union of the given theories, which satisfy the Integrity Constraints. It is defined as:
$Comb_3(\{T_1, \cdots, T_k\}, IC) \stackrel{\text{def}}{=} \{X : X \in MAXCONS(T) \text{ and } X \cup IC \text{ is consistent.}\}$

Considering Example 4.1 we have $MAXCONS(T) = \{T_1 = \{a, b, c\}, T_2 = \{a, c, \neg a \vee \neg b\}, T_3 = \{b, c, \neg a \vee \neg b\}\}$
Of the three theories in $MAXCONS(T)$ only T_3 is consistent with respect to $IC = \{\neg a \vee \neg c\}$. Hence, the third approach considers only T_3.

Theorem 9 $Comb_3$ is a combination function. □

4.1 A More Practical Approach

In all the previous combination functions we searched for all the maximally consistent subsets of the inconsistent theories. This is an extremely time consuming process since it requires consistency checks in each step of the algorithm (Algorithm 3.2). In this section we would like to present a more practical algorithm.

In addition we are motivated by the following argument. We are combining several theories. Each of them is consistent and the inconsistency arose from the union of the theories and we do not know which of the theories being combined contains erroneous information. We would like to search for a combination that will include as much information as possible from the original theories. By "as much", we mean as many clauses as possible, i.e. theories that are maximal w.r.t. cardinality rather than with respect to inclusion.

The following $MAXCONS3$ serves both purposes.

Algorithm 4.1 Procedure $MAXCONS3(P, IC)$

```
MAXCON = ∅.
CHECK = {P}.
Found = false
While not Found
   begin
   TEMP = ∅.
   For all Pi ∈ CHECK do
      Begin
      If Pi ∪ IC is consistent then
         Begin
         Found=True.
         MAXCON = MAXCON ∪ {Pi ∪ IC}.
         End
      else
         begin
         {** Let Pi be the set of sentences {Ci1, ···, Cin}. **}
         For j = 1···n do TEMP = TEMP ∪ {Pi − {Cij}} od
         End
      End    CHECK = TEMP.
   End
```

$MAXCONS3(P, IC)$:= *maximal elements of* $MAXCON$.
end

Definition 4.8 $Comb_4(\{T_1, \cdots, T_k\}, IC) \stackrel{\text{def}}{=} MAXCONS3(T_1 \cup \cdots \cup T_k, IC)$.

Theorem 10 $Comb_4(\{T_1, \cdots, T_k\}, IC) =$ maximal sets WRT cardinality of $Comb_1(\{T_1, \cdots, T_k\}, IC)$.

The following example illustrates the difference between this approach and the $MAXCONS(P)$ approach.

Example 4.2 Consider the theory P below:

d
$\neg d$
$d \rightarrow e$
$d \rightarrow \neg e$

$MAXCONS3(P)$ contains exactly one element, viz. $\{\neg d, d \rightarrow e, d \rightarrow \neg e\}$. Note, however, that in addition to this set, $MAXCONS(P)$ contains $\{d, d \rightarrow e\}$, $\{d, d \rightarrow \neg e\}$, $\{\neg d, d \rightarrow \neg e\}$, $\{d, d \rightarrow \neg e\}$, $\{d, d \rightarrow e\}$, $\{\neg d, d \rightarrow e\}$.

The following result is easy to establish.

Corollary 4.2 Let P be any first order theory. $MAXCONS3(P) \subseteq MAXCON(P)$. □

We note that the inferences of $MAXCONS3$ is less cautious and more optimistic than those $MAXCONS$, i.e. everything that is infered from $MAXCONS$ is also inferred from $MAXCONS3$, but there is more information that is inferred from $MAXCONS3$ than from $MAXCONS$.

4.2 Relationship with Updating of Theories

The problem of updating theories and revising beliefs has been extensively studied [FUV83, FKUV86, Gar88, Dal88, KM89, GM88, Sat88]. In brief, an update may take one of two forms:

1. An Insertion: In this, some *new* information is added to T.

2. A Deletion: Here, either a formula F in T is deleted, or a formula entailed by T is deleted.

Inserting a formula F into T may lead to an inconsistency. In this case, the *result* of the insertion must be defined in such a way that the inconsistency is properly handled. Deletions do not give rise to inconsistencies (unless T was already inconsistent).

There is some resemblance between insertions and the combination of theories we have discussed earlier. Insertions into theories may be viewed as a special case of our framework: take the sentence F to be inserted to be an integrity constraint, i.e. $IC = \{F\}$ and now compute $MAXCONS(T, IC)$. This is the gist of Theorem 12 below.

One may wonder whether *combining* theory T_1 with theory T_2 may be accomplished by inserting each element of T_2 into T_1. This is not true in general (cf. Example 4.3 below). The reason for this is that we do not have any priorities over the set of the combined theories.[1] For example when we insert elements of T_2 into T_1 one by one, then the last element of T_2 to be inserted is accorded the status of an integrity constraint, even though this last element of T_2 may not be present in all maximal consistent subsets of $T_1 \cup T_2$.

Fagin, Ullman and Vardi [FUV83] present a theory of updating theories. Before discussing the relationship between updating and combining we discuss the theory updating approach of Fagin, Ullman and Vardi [FUV83].

Definition 4.9 [FUV83] Let T be a theory and $T*$ be the set of clauses logically implied by T. A theory S is said to accomplish the *insertion* of a clause σ into T if $\sigma \in S$. A theory S is said to accomplish the *deletion* of a clause σ into T if $\sigma \notin S*$.

Definition 4.10 [FUV83] Let T, T_1 and T_2 be theories. T_1 has *fewer insertions* than T_2, if $T_1 - T \subset T_2 - T$. T_1 has *no more insertions* than T_2, with respect to T, if $T_1 - T \subseteq T_2 - T$. T_1 has the *same insertions* as T_2, with respect to T, if $T_1 - T = T_2 - T$ Let T, T_1 and T_2 be theories. T_1 has *fewer deletions* than T_2, if $T - T_1 \subset T - T_2$. T_1 has *no more deletions* than T_2, with respect to T, if $T - T_1 \subseteq T - T_2$. T_1 has the *same deletions* as T_2, with respect to T, if $T - T_1 \subseteq T - T_2$.

[1] Using Gärdenfors' terminology, our combination functions do not satisfy axiom K_2^+ ([Gar88] page 48)

Definition 4.11 [FUV83] Let T, T_1 and T_2 be theories. T_1 accomplishes an update u (could be an insertion or a deletion) of T *with a smaller change* than T_2 if both T_1 and T_2 accomplish u, and either T_1 has fewer deletions than T_2 or T_1 has the same deletions as T_2 but T_1 has fewer insertions than T_2.

Definition 4.12 [FUV83] A theory S accomplishes an update u of a theory T *minimally* if there is no theory S' that accomplishes u with a smaller change than S.

Theorem 11 [FUV83] Let S and T be theories and let σ be a sentence. Then,
1. S accomplishes the deletion of σ from T minimally if and only if S is a maximal subset of T that is consistent with $\neg\sigma$, and
2. S accomplishes the insertion of σ from T minimally if and only if $S \cap T$ is a maximal subset of T that is consistent with σ. □

The following theorem describes a relationship between combining theories and updating theories, when the union of the theories to be combined is consistent.

Theorem 12 Let $T_1, \cdots, T_k$ be theories to be combined in the presence of a finite set IC of integrity constraints. Let σ_{ic} be the conjunction of the integrity constraints in IC. If $T_1 \cup \cdots \cup T_k$ is consistent then, $Comb_1(\{T_1, \cdots, T_k\}, \{\sigma_{ic}\}) = \{\ X | X$ accomplishes the insertion of σ_{ic} into $T_1 \cup \cdots \cup T_k$ minimally. $\}$ □

Proof: $[Comb_1(\{T_1, \cdots, T_k\}, \sigma_{ic}) \subseteq \{X | X$ accomplishes the insertion of σ_{ic} into $T_1 \cup \cdots \cup T_k$ minimally $\}]$ Suppose $X \in Comb_1(\{T_1, \cdots, T_k\}, \sigma_{ic})$. Then X is a maximally consistent set such that $\{\sigma_{ic}\} \subseteq X \subseteq T_1 \cup \cdots \cup T_k \cup \{\sigma_{ic}\}$. Hence, $\sigma_{ic} \in X$. It now suffices to show that X is a maximally consistent subset of $T_1 \cup \cdots T_k \cup \{\sigma_{ic}\}$. But this is true. Hence $X \in \{X | X$ accomplishes the insertion of σ_{ic} into $T_1 \cup \cdots \cup T_k$ minimally $\}$.

$[\{X | X$ accomplishes the insertion of σ_{ic} into $T_1 \cup \cdots \cup T_k$ minimally $\} \subseteq Comb_1(\{T_1, \cdots, T_k\}, \sigma_{ic})]$
The proof is similar. □

Theorem 12 above demonstrates that the Fagin et. al. [FKUV86] framework for inserting sentences into theories may be captured in our framework. The example below shows that successive insertions of sentences of theory T_2 into theory T_1 does not correctly capture the combination of theories.

Example 4.3 Suppose $T_1 = \{$
$a \rightarrow c,$
$a \rightarrow \neg b,$
$a \rightarrow b\},$

and $T_2 = \{a\}$. Inserting a into T_1 gives us two sets that accomplish this insertion: $S_1 = \{\ a$
$a \rightarrow c$
$a \rightarrow \neg b\}$

and $S_2 = \{\ a$
$a \rightarrow c$
$a \rightarrow b\}$

c is true in both S_1 and S_2. Note however, that according to $MAXCONS(T_1 \cup T_2)$, c is not true in T_1 which is a maximally consistent subset of $T_1 \cup T_2$.

We now the use of the theory developed by Fagin et al. [FKUV86] for updates in flocks to compare with our problem of combining theories.

Definition 4.13 [FKUV86] Let $\mathcal{S} = \{S_1, \cdots, S_n\}$ be a flock. A flock $\mathcal{T} = \{T_1, \cdots, T_n\}$ accomplishes an update u of $\mathcal{S}$ minimally if T_i accomplishes the update of S_i minimally. □

Definition 4.14 [FKUV86] Let $\mathcal{S}$ be a flock and $\mathcal{S}_1, \cdots, \mathcal{S}_k$ be the flocks that accomplish an update u of $\mathcal{S}$ minimally. Then the result of u is the flock $\bigcup_{1 \leq i \leq k} \mathcal{S}_i$.
□

We would like to change Definition 4.14 so that the resulting flock consists of maximal elements only. Formally,

Definition 4.15 Let $\mathcal{S}$ be a flock and $\mathcal{S}_1, \cdots, \mathcal{S}_k$ be the flocks that accomplish an update u of $\mathcal{S}$ minimally. Then the result of u is the flock consisting maximal elements of $\bigcup_{1 \leq i \leq k} \mathcal{S}_i$. □

The following result is immediate.

Theorem 13 Let $T_1, \cdots, T_k$ be a set of theories and let IC be a finite set of integrity constraints. Let $\mathcal{T}$ be the flock $MAXCONS(T_1 \cup \cdots \cup T_k)$. Let σ_{ic} be the conjunction of the integrity constraints in IC. $Comb_1(\{T_1, \cdots, T_k\}, \{\sigma_{ic}\})$ = the flock obtained by updating $\mathcal{T}$ with σ_{ic} by using Definition 4.15. □

5 Combining Prioritized Theories

The different knowledge bases that have to be combined might have different priorities associated with them. Intuitively, there might be compelling reasons that cause one knowledge base to be preferable to another. In such a case we would like to use this priority information while combining the knowledge bases. For example, if a knowledge base with higher priority or believability directly contradicts another knowledge base with a lower priority, with respect to a certain aspect we might want the combined theory to contain the point of view of the knowledge base with the higher priority. In this section we formalize what we mean by combining prioritized theories.

In the context of the murder example of Section 2, the information provided by Bill (the younger witness who was also closer to the scene of action) may seem more credible to the police who may then pursue further investigations based on his version of events as opposed to John's version.

As a first step suppose we have theories $T_1, \cdots, T_k$ to be combined with the priority relation (a total order) where the priority of T_i is less than priority of T_j iff $i < j$. With a slight abuse of notation it can be written as, $T_1 \prec T_2 \prec \cdots \prec T_k$. As always the integrity constraints have the highest priority, i.e. $T_k \prec IC$. There are two distinct approaches to combine the theories that come to mind immediately; ***bottom-up*** and ***top-down***.

In the bottom-up approach we start by combining T_1 and T_2 with preference to T_2. The combined theory is then a set of theories defined as $\{T : T_2 \subseteq T$ and $T - T_2$ is a maximal subset of T_1 such that T is consistent$\}$. The result is then combined with T_3 with preference to T_3. This means that each theory in the result is combined with T_3 and the final result is the set which is the union of particular results (each result is a set). This continues until T_k. The result is then combined with IC with preference to IC.

Algorithm 5.1 Procedure $Comb-botup(Comb_i, T_1 \preceq \cdots \preceq T_k, IC)$

$T_{k+1} = IC$

$Comb = \{T_1\}$

For $i = 1$ *to* k *do*

begin

$Temp = \emptyset$

For all T *in* $Comb$ *do*

$Temp = Temp \cup Comb_i(\{T\}, T_{i+1})$

$Comb =$ *maximal elements of* $Temp$

end

Output(Comb)

In the top-down approach we start with combining T_k and IC with preference to IC. The result is combined with T_{k-1} with preference to the theories in the result. This is continued until T_1.

Algorithm 5.2 Procedure $Comb-topdn(Comb_i, T_1 \preceq \cdots \preceq T_k, IC)$

$T_{k+1} = IC$

$Comb = \{T_{k+1}\}$

For $i = k$ *to* 1 *do*

begin

$Temp = \emptyset$

For all T *in* $Comb$ *do*

$Temp = Temp \cup Comb_i(\{T_i\}, T)$

$Comb =$ *maximal elements of* $Temp$

end

Output(Comb)

Example 5.1 Consider, $T_3 = \{\neg d\}$, $T_2 = \{c \rightarrow \neg d; d\}$, $T_1 = \{c\}$, and $IC = \emptyset$, with the priorities $IC \prec T_3 \prec T_2 \prec T_1$.

Top-down:

Combining T_3 and T_2 we obtain $T_{32} = \{\{c \rightarrow \neg d; \neg d\}\}$.

Combining T_{32} with T_1 we obtain $T_{321} = \{\{c \rightarrow \neg d; \neg d; c\}\}$.

Bottom-up:

Combining T_1 and T_2 we obtain $T_{12} = \{c \rightarrow \neg d; d\}$, Combining T_{12} with T_3 we obtain $T_{123} = \{c \rightarrow \neg d; \neg d\}$,

Hence in this case the result obtained by the top-down approach is different from the result obtained by the bottom-up approach. We note also, that even though in our case there is only one maximal consistent set in each of the combinations, it is not the case in general.

Here, again, one may wonder whether the *combining* of prioritized theories is similar to the updating problem. Since the motivations for the two problems are different so are the results. The updating problem can be simulated by the *bottom-up* procedure for combining *prioritized* theories which is different from both the *top-down* procedure and the procedure for combining non-prioritized theories. □

While speaking of the combination of prioritized theories, we briefly mention that it is possible that we have *prioritized groups* of theories $G_1, \ldots, G_n$. Any two theories in the same group have the same priority; however, the groups themselves have priorities $G_1 \prec G_2 \prec \cdots \prec G_n$. The most straightforward way of combining the resulting multitude of theories is to proceed as follows:

Step 1: For all $1 \le i \le n$, set $S_i = MAXCON(\bigcup_{T \in G_i} T)$. Thus, at this stage, for any $1 \le i \le n$, all theories in the group G_i have been combined together. Each S_i is thus a set of theories.

Step 2. Construct $S = (S_1 \times \cdots \times S_n)$, i.e. S is the Cartesian product of the S_i's.

Step 3. Let $S = \{V_1, \ldots, V_k\}$ where each V_i is of the form $\langle s_1, \ldots, s_n \rangle$ where $s_j \in S_j$ for all $1 \le j \le n$.

Step 4. Combine theories $s_1, \ldots,_n$ with priorities $s_1 \prec s_2 \prec \cdots \prec s_n$. Do this for all $V_i \in S$. Let the resulting theories be $\{TH_1 \ldots, TH_k\}$.

Step 5. Choose the maximal elements of $\{TH_1 \ldots, TH_k\}$.

A detailed discussion about combining prioritized groups of theories is beyond the scope of this paper.

In order to see how the work described thus far relates to existing work on reasoning with inconsistency, consider the following program P:

p

$\neg p$

$p \rightarrow q$

$\neg p \rightarrow q$

$p \,\&\, \neg p \rightarrow r$

Using the annotated logic semantics of Blair and Subrahmanian [BS88, BS89] and later improved by Kifer and Lozinskii [KL89], it would be possible to infer r even though r depends on a somewhat shaky justification, viz. $(p \,\&\, \neg p)$. However, the annotated logic semantics does not allow us to conclude $\neg r$ or $\neg q$. Clearly, the fact that $\neg r$ and $\neg q$ cannot be inferred corresponds quite well with our intuition. The $MAXCONS(P)$ approach would allow us to conclude r. At the same time, neither $\neg r$ nor $\neg q$ can be concluded.

Considering the same program P, the semantics of Gelfond and Lifschitz [GL] and Kowalski and Sadri [KS] allows us to conclude everything. In particular, r can be concluded (in the same way as in [BS88, BS89]), but in addition, both $\neg r$ and $\neg q$ may be concluded.

One advantage of the Blair and Subrahmanian [BS88, BS89, KL89] approach is that in the case of programs containing function symbols, their semantics leads to a semi-decidable consequence relation. When function symbols are present, this does not appear to be true for the $MAXCONS(P)$ semantics.

6 Conclusions

In this paper, we have developed formal techniques for combining multiple knowledge bases. This is directly related to the construction of expert systems because knowledge contributed by different experts may be mutually conflicting. This may be due to carelessness or perhaps a lack of knowledge or, more seriously, due to a genuine disagreement amongst the experts.

Based on the *cautious* semantics suggested by Grant and Subrahmanian [GS90], we have developed various ways of reasoning about combinations of theories in the presence of integrity constraints. We have argued that these different methods are desirable under different circumstances. It is up to the individual researcher to decide which semantics is most appropriate for his/her use.

Acknowledgements. We wish to express our appreciation to the National Science Foundation for their support of our work under grant number IRI-86-09170 and the Army Research Office under grant number DAAG-29-85-K-0-177. Sarit Kraus

was partially supported by National Science Foundation grant IRI-8907122.

References

[BKM89] C. Baral, S. Kraus, and J. Minker. Combining Multiple Knowledge Bases. Technical Report UMIACS TR 89-90, CS TR 2316, Department of Computer Science and UMIACS, University of Maryland, College Park, Md 20742, September 1989. To appear in *IEEE Transactions on Data and Knowledge Engineering.*

[BS88] H. A. Blair and V.S. Subrahmanian. Paraconsistent Foundations for Logic Programming. *Journal of Non-Classical Logic*, 5,2, 1988.

[BS89] H. A. Blair and V.S. Subrahmanian. Paraconsistent Logic Programming. *Theoretical Computer Science*, 68:135–154, 1989.

[Dal88] M. Dalal. Investigations into a Theory of Knowledge Base Revision: Preliminary Report. In *AAAI*, pages 475–479, 1988.

[dCSV89] N.C.A. da Costa, V.S. Subrahmanian, and C. Vago. The Paraconsistent Logics $\mathcal{P}\mathbf{T}$. *to appear in Zeitschrift fur Mathematische Logik und Grundlagen der Mathematik*, 37, 1991.

[FKUV86] R. Fagin, G. Kuper, J. Ullman, and M. Vardi. Updating Logical Databases. In *Advances in Computing Research*, volume 3, pages 1–18, 1986.

[FUV83] R. Fagin, J.D. Ullman, and M.Y. Vardi. On the Semantics of Updates in Databases. In *ACM SIGACT/SIGMOD Symposium on Principles of Database Systems*, pages 352–365, 1983.

[Gar88] P. Gardenford. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, 1988.

[GL] M. Gelfond and V. Lifschitz. Logic Programs with Classical Negation. In *Proc. 7th Intl. Conf. on Logic Programming*, pages 579–597, MIT Press.

[GM88] P. Gardenfors and David Makinson. Revisions of Knowledge Systems using Epistemic Entrenchment. In *Proc. 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, 1988.

[Gra74] J. Grant. Incomplete Models. *Notre Dame J. of Formal Logic*, XV(4):601–607, 1974.

[Gra75] J. Grant. Inconsistent and Incomplete Logics. *Mathematics Magazine of Mathematical Association of America*, 48(3):154–159, 1975.

[Gra77] J. Grant. Incompleteness and Inconsistency in Propositional Logic. *Relevent Logic Newsletter*, 1977.

[Gra78] J. Grant. Classifications for Inconsistent Theories. *Notre Dame J. of Formal Logic*, XIX(3):435–444, 1978.

[GS90] J. Grant and V. S. Subrahmanian. Reasoning in Inconsistent Databases, 1990. manuscript.

[KM89] H. Katsuno and A. Mendelzon. A Unified View of Propositional Knowledge Base Updates. In *Eleventh International Joint Conference on AI*, Detroit, August 89.

[KL89] M. Kifer and E. Lozinskii. A Logic for Reasoning with Inconsistency. to appear in *J. of Automated Reasoning.* Preliminary version in: Proc. Logic in Computer Science 1989.

[KS] R. Kowalski and F. Sadri. Logic Programs with Exceptions. In *Proc. 7th Intl. Conf. on Logic Programming*, pages 598–613, MIT Press.

[Sat88] Ken Satoh. Non-Monotonic Reasoning by Belief Revision. In *Proc. of the international Conference on Fifth Generation Computer Systems*, pages 455–462, 1988.

[Sub89] V. S. Subrahmanian. Paraconsistent Disjunctive Deductive Databases. to appear in *Theoretical Computer Science*, 1992.